JOINT RECONSTRUCTION OF NOISY HIGH-RESOLUTION MR IMAGE SEQUENCES

Justin P. Haldar and Zhi-Pei Liang

Department of Electrical and Computer Engineering, Beckman Institute, University of Illinois at Urbana-Champaign

ABSTRACT

Quantitative MR studies often utilize sequences of coregistered images, where the contrast in each image frame is experimentally manipulated to enable the regression of important physical parameters. However, the potential of these experiments has been limited for high-resolution biological studies because of long acquisition times and limited signal-to-noise ratio. This work presents a new approach for the reconstruction of an image sequence from noisy data, using a statistical model that incorporates an implicit line-site prior to take advantage of the high level of inter-frame correlation between spatial image features. Reconstructions are efficiently computed using a globally-convergent half-quadratic iterative algorithm, and the proposed optimization procedure enables precise characterization of resolution and noise properties.

Index Terms— magnetic resonance imaging, image reconstruction, denoising, half-quadratic regularization, image sequences

1. INTRODUCTION

Contrast in MR images is dependent on both the physical properties of the imaging subject and on the experimental parameters used to acquire data. The goal of many quantitative MR experiments is to directly estimate these physical characteristics, since this higher-level information can provide direct insight into the structure, function, and viability of biological tissues.

In practice, most quantitative MR experiments (e.g., diffusion imaging, perfusion imaging, and relaxometry) acquire a sequence of coregistered images, where the contrast of each image frame is carefully manipulated. The contrast variations observed in each voxel of the image sequence are then fit to an appropriate physical model of the contrast mechanism, finally resulting in quantitative spatial maps of the parameters of interest.

In this work, we consider the case where a length-Q sequence of coregistered images is acquired with different contrasts from a static object. If we let $d_q(\mathbf{k}_m)$ represent the acquired *k*-space data for the qth image frame, and let $\rho_q(\mathbf{x})$ represent the corresponding image, then the acquired experimental data can be modeled as

$$d_q(\mathbf{k}_m) = \int \rho_q(\mathbf{x}) \exp(-\imath 2\pi \mathbf{k}_m \cdot \mathbf{x})\, d\mathbf{x} + \eta_{mq}, \quad m \in \{0, 1, \ldots, M-1\}, \;\; q \in \{1, 2, \ldots, Q\} \qquad (1)$$

where M is the total number of spatial encodings and η_{mq} is complex white Gaussian noise with variance σ^2.

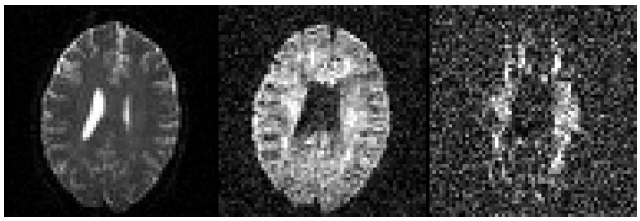

Fig. 1. Three frames from a DSI image sequence. Due to the large number of frames required for DSI (i.e., $Q = 515$), full DSI acquisitions can take more than 25 minutes, and the resulting data still suffers from relatively low resolution and SNR. (Data courtesy of Dr. V. J. Wedeen of the MGH Martinos Center for Biomedical Imaging and Harvard Medical School)

There are two important practical limitations for these imaging schemes. First, the imaging time is inherently long due to the number of samples that must be collected to sufficiently encode the physical properties of interest with sufficient spatial resolution. Second, depending on the specific contrast mechanisms utilized by the experiment, the data can have very low signal-to-noise ratio (SNR). This is especially problematic, because it further limits the achievable imaging speed. As a result, many quantitative *in vivo* MR experiments use relatively large voxel sizes (i.e., 3-5 mm or more along each dimension), and still suffer from limited SNR. Figure 1 illustrates these effects using a typical example from human brain diffusion spectrum imaging (DSI) [1].

This paper proposes a new method to address the SNR problem. The method utilizes a specially-designed statistical reconstruction that performs joint estimation of the entire image sequence, and exploits the shared structure inherent to these kinds of experiments.

The work reported in this paper was supported in part by the NSF GRFP, and by research grants NIH-P41-EB03631-16 and NIH-R01-CA098717.

2. PROPOSED METHOD

2.1. Joint Feature-Preserving Statistical Reconstruction

We propose to reconstruct the set of image frames using a quasi-Bayesian optimality criterion

$$\left\{\hat{\boldsymbol{\rho}}_1, \hat{\boldsymbol{\rho}}_2, \ldots, \hat{\boldsymbol{\rho}}_Q\right\} = \underset{\{\boldsymbol{\rho}_1, \boldsymbol{\rho}_2, \ldots, \boldsymbol{\rho}_Q\}}{\arg\min} \left\{ \begin{array}{c} \left[\sum_{q=1}^{Q} w_q^2 \left\|\mathbf{F}\boldsymbol{\rho}_q - \mathbf{d}_q\right\|_2^2\right] \\ +\lambda R\left(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2, \ldots, \boldsymbol{\rho}_Q\right) \end{array} \right\}, \quad (2)$$

where $\mathbf{F}$ represents the Fourier imaging operator, $\boldsymbol{\rho}_q$ is the length-N vector of voxel coefficients for the qth image frame, $\mathbf{d}_q$ is the corresponding length-M vector of data samples, the w_q are weighting parameters which will be discussed later in this section, λ is a regularization parameter, and $R(\cdot)$ is a regularization functional. In this framework, choosing $R(\cdot)$ to appropriately incorporate prior information is essential, and this paper determines $R(\cdot)$ based on the following considerations:

- Images from typical MRI experiments usually consist of "smooth" regions separated by edges. Smoothness constraints can be used to reduce the effects of noise, although edges need to be preserved to avoid amplifying partial volume effects.
- The edge structures seen in different frames of an image sequence have strong correlation. For example, object support boundaries will exist in every image, regardless of the image contrast. Edge structures should be imposed in a joint fashion.
- Resolution and SNR properties of the resulting reconstructions should be characterizeable. Moreover, it is important that different reconstructed image frames have identical resolution properties, so that the same voxel in different image frames accurately reflects signal from the same spatial spin population. This is necessary to ensure the validity of any parametric voxel-by-voxel model-fitting procedure that would follow reconstruction of the image sequence.
- The resulting reconstruction should be relatively insensitive to the choice of regularization parameters and initialization.

Based on the above factors, we choose a joint smoothness prior of the form

$$R\left(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2, \ldots, \boldsymbol{\rho}_Q\right) = \sum_{n=1}^{N} \sum_{\substack{m \in \Omega_n \\ m > n}} \Psi\left(\sqrt{\sum_{q=1}^{Q} w_q^2 \left|\rho_{q,n} - \rho_{q,m}\right|^2}\right), \quad (3)$$

where Ω_n is the set of all voxels adjacent to the nth voxel, and $\Psi(\cdot)$ is an appropriate robust-statistical cost function. It can be shown (e.g., following derivations in [2, 3]) that this choice is equivalent to imposing a quasi-Bayesian line-site prior on the image frames. The line-site prior, first proposed by Geman and Geman [4], provides a natural way to model the edge structure within images. It does this by imposing that spatially-adjacent voxels are related to each other through new variables called *line-sites*, which are used to explicitly model this structure. The quasi-Bayesian restoration problem of Eq. 2 is then equivalent to jointly estimating the image and its edge structure, where the edge modeling allows for the preservation of important image features.

In contrast to reconstruction algorithms which impose a line-site prior to each frame independently, Eq. 3 results in a line-site prior that is shared between all image frames. The proof of this follows from the derivations in [2, 3], and the shared edge map will appear explicitly in the proposed optimization scheme. The use of a shared edge map leverages the correlation between different frames to improve the quality of the jointly-reconstructed images.

In what follows, we shall consider $\Psi(\cdot)$ functions of the form

$$\Psi(t) = \begin{cases} \frac{2}{\beta}\left(t^2+\epsilon\right)^{\frac{\beta}{2}} - \frac{2\epsilon^{\frac{\beta}{2}}}{\beta}, & \text{when } 0 \le t < \alpha \\ 2\left(\alpha^2+\epsilon\right)^{\frac{\beta-1}{2}}\sqrt{t^2+\epsilon} - \frac{2\epsilon^{\frac{\beta}{2}}}{\beta} & \\ \quad -2\left(\alpha^2+\epsilon\right)^{\frac{\beta-1}{2}}\left(1-\frac{1}{\beta}\right)\sqrt{\alpha^2+\epsilon}, & \\ \quad \text{when } t \ge \alpha & \end{cases} \quad (4)$$

where ϵ is a small positive number (we use $\epsilon = 10^{-10}$), β is a number such that $1 \le \beta < 2$ (we use $\beta = 1.99$), and α is nonnegative. This cost function may look complicated, but is actually proportional to a simple approximation of Huber's minimax robust estimator, which was modified to ensure convergence of our proposed algorithm for optimizing Eq. 2 and to simplify notation in future equations. Huber's minimax estimator has been used previously in image processing (see [3] and its references for further discussion). The parameter α plays a special role in this situation; essentially, the magnitude of an edge has to be stronger than α to be declared as an edge. Therefore, using a larger α improves robustness against declaring false edges, though also makes it easier to not identify actual edge structures. Note that with $\alpha = 0$ and $Q = 1$, Eq. 3 is the penalty functional associated with a common formulation of the ubiquitous Total Variation regularization problem.

2.2. Optimization Algorithm

With the robust cost function of Eq. 4, the following half-quadratic optimization procedure is guaranteed to converge to the unique optimal solution from any initialization [5–7], as long as the acquisition includes $\mathbf{k} = \mathbf{0}$, the center of *k*-space.

1. As a first step, set the estimated image sequence equal to an initial guess $\left\{\hat{\boldsymbol{\rho}}_1^{(0)}, \hat{\boldsymbol{\rho}}_2^{(0)}, \ldots, \hat{\boldsymbol{\rho}}_Q^{(0)}\right\}$ (e.g., set all voxel values based on an initial noisy Fourier reconstruction). A good initial guess is not essential, because the proposed algorithm converges to the unique optimal solution from any starting point.

2. At the jth iteration, define line-site variables for each voxel pair as

$$\ell_{n,m}^{(j)} = \begin{cases} \frac{1}{\left(\left(t_{n,m}^{(j)}\right)^2+\epsilon\right)^{1-\frac{\beta}{2}}}, & 0 \le t_{n,m}^{(j)} < \alpha \\ \frac{\left(\alpha^2+\epsilon\right)^{\frac{\beta-1}{2}}}{\sqrt{\left(t_{n,m}^{(j)}\right)^2+\epsilon}}, & t_{n,m}^{(j)} \ge \alpha \end{cases} \quad (5)$$

where $t_{n,m}^{(j)}$ is given by

$$t_{n,m}^{(j)} = \sqrt{\sum_{q=1}^{Q} w_q^2 \left|\rho_{q,n}^{(j)} - \rho_{q,m}^{(j)}\right|^2}. \quad (6)$$

3. Update each image frame according to

$$\hat{\boldsymbol{\rho}}_q^{(j+1)} = \arg\min_{\boldsymbol{\rho}_q} \left\|\mathbf{F}\boldsymbol{\rho}_q - \mathbf{d}_q\right\|_2^2 + \lambda R^{(j)}\left(\boldsymbol{\rho}_q\right), \quad (7)$$

with

$$R^{(j)}\left(\boldsymbol{\rho}_q\right) = \sum_{n=1}^{N} \sum_{\substack{m \in \Omega_n \\ m>n}} \ell_{n,m}^{(j)} \left|\rho_{q,n} - \rho_{q,m}\right|^2. \quad (8)$$

Note that this is a purely quadratic optimization problem, and has the closed form solution

$$\hat{\boldsymbol{\rho}}_q^{(j+1)} = \left(\mathbf{F}^H\mathbf{F} + \lambda\mathbf{D}^H\mathbf{W}^{(j)}\mathbf{D}\right)^{-1}\mathbf{F}^H\mathbf{d}_q, \quad (9)$$

where $\mathbf{D}$ is a sparse matrix such that $\mathbf{D}\boldsymbol{\rho}_q$ computes finite differences of the form $(\rho_{q,n} - \rho_{q,m})$, and $\mathbf{W}^{(j)}$ is a diagonal matrix with diagonal elements equal to the line-site variable $\ell_{n,m}^{(j)}$ for the corresponding finite difference computed by $\mathbf{D}$. While all the matrices in Eq. 9 are very large, they are also highly structured, so the matrix inversion can be performed efficiently using standard iterative methods like the conjugate-gradient algorithm.

4. Repeat steps 2 and 3 until the iterations converge, i.e., when $\left\|\hat{\boldsymbol{\rho}}_q^{(j+1)} - \hat{\boldsymbol{\rho}}_q^{(j)}\right\|_2^2$ is negligible for each image frame. At convergence, the reconstructed images will be the unique optimal solution to the original quasi-Bayesian problem.

Note that the iterative procedure uses line-sites, and that the value of the line-site depends on a weighted average edge-strength seen in all images through Eq. 6, where the weighting is defined through the w_q parameters. Setting $w_q = 1$ for all images can be problematic, as it will give unequal weight to images which are scaled differently from one another. For example, in DSI, images with very light diffusion weighting typically have much larger intensities relative to images with high amounts of diffusion weighting. In such a case and with $w_q = 1$, the estimated edge-maps will be dominated by structures visible in the images with low amounts of diffusion weighting. As can be seen in Fig. 1, this type of weighting will neglect some of the higher-level structures that only become apparent with high levels of diffusion weighting. We generally choose the w_q parameters inversely proportional to the average magnitude of the signal within the parenchyma from a noisy reconstruction of the corresponding image frame.

2.3. Properties

The proposed method has several desirable properties. First, as illustrated by Eq. 9, for a given estimated edge map, the estimated image $\hat{\rho}_q(\mathbf{x})$ is linear with respect to the data $d_q(\mathbf{k}_m)$. This linearity provides analytic methods to analyze the resolution and noise properties of the reconstructed image, similar to what is done in [8]. In particular, we can define a *spatial response function* (SRF) for each voxel, which is similar to the point spread function typically used to characterize resolution in conventional image reconstruction. Specifically, the SRF for the nth voxel is given by

$$h_n(\mathbf{x}) = \sum_{m=1}^{M} \mathbf{G}_{nm} e^{-i2\pi\mathbf{k}_m\cdot\mathbf{x}}, \quad (10)$$

where $\mathbf{G}_{nm}$ are entries from the reconstruction matrix in Eq. 9. It is easy to show that the value of the nth reconstructed voxel coefficient is given by

$$\hat{\rho}_{q,n} = \int \rho_q(\mathbf{x})\, h_n(\mathbf{x})\, d\mathbf{x} + \bar{\eta}_n, \quad (11)$$

where $\bar{\eta}_n$ is a zero-mean Gaussian random variable with variance

$$\sigma^2 \sum_m \left|\mathbf{G}_{nm}\right|^2. \quad (12)$$

It is important to note that while the SRF may be position dependent, the proposed method guarantees that the different image frames share the same SRFs, and thus have identical resolution and noise properties. Thus, the spin population represented by a given voxel is consistent for every image. In addition, unlike approximate characterizations using the *local impulse response* [9], Eqs. 10 and 12 are valid even if the iterative procedure is halted prior to convergence.

Second, analysis of SRFs shows that reconstruction including an edge map has a number of nice advantages. As expected, both the noise variance and resolution of reconstructed images decreased as stronger emphasis is placed on image smoothness by increasing λ. This is the same type of effect typically seen when apodization is applied in standard Fourier reconstruction. However, by using an edge map, we can prevent signal from leaking across a known edge, and can thus avoid exacerbating partial volume effects, ensuring that local signal averaging is only performed in regions that should be relatively homogeneous in the first place. The regularization parameter λ can be chosen based on the desired trade-off between resolution and SNR.

Third, the reconstruction results are not sensitive to the initialization of the algorithm, and are less sensitive to the choice of reconstruction parameters relative to nonconvex approaches. It is also worth pointing out that the proposed algorithm reconstructs images directly from the acquired *k*-space data, which allows us to accomodate non-Cartesian *k*-space trajectories and parallel imaging in a statistically optimal way, and avoid the spatial noise correlation and ringing artifacts resulting from zero-padding or non-Cartesian reconstruction.

3. RESULTS

Figures 2 and 3 show the results of the proposed algorithm applied to two different diffusion imaging experiments. While the resolution in smooth regions of the image is slightly reduced, high-quality reconstructions are achievable in a fraction of the previously-required imaging time.

4. CONCLUSION

This paper presents a novel method for the joint reconstruction of different frames in an image sequence. The formulation uses the link between robust statistics and half-quadratic regularization to develop a quasi-Bayesian prior that is well-suited for the preservation of joint image features. The proposed optimization procedure is globally-convergent, and enables precise characterization of the final resolution and noise properties.

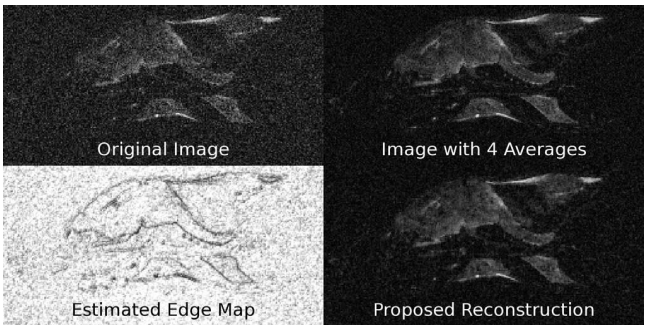

Fig. 2. One image frame from a 12-frame diffusion-weighted mouse-brain sequence. The proposed method improves image SNR, while preserving important structural features.

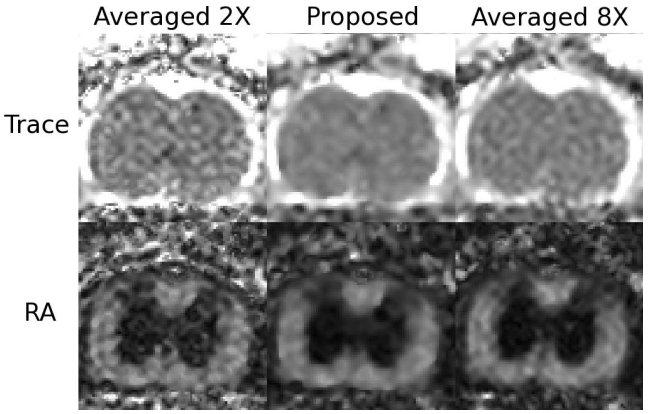

Fig. 3. Applying the proposed algorithm applied to mouse spine diffusion tensor imaging data (courtesy of Drs. J. H. Kim and S.-K. Song of the Washington University in St. Louis) results in reduced noise bias and Gibbs phenomena in estimated parametric maps of important diffusion parameters.

5. REFERENCES

[1] V. J. Wedeen, P. Hagmann, W.-Y. I. Tseng, T. G. Reese, and R. M. Weisskoff, "Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging," *Magn Reson Med*, vol. 54, pp. 1377–1386, 2005.

[2] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Trans Patt Anal Mach Int*, vol. 14, pp. 367–383, 1992.

[3] M. J. Black and A. Rangarajan, "On the unification of line processes, outlier rejection, and robust statistics with applications in early vision," *Int J Comput Vis*, vol. 19, pp. 57–91, 1996.

[4] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images," *IEEE Trans Patt Anal Mach Int*, vol. 6, pp. 721–741, 1984.

[5] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans Imag Proc*, vol. 6, pp. 298–310, 1997.

[6] A. H. Delaney and Y. Bresler, "Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography," *IEEE Trans Imag Proc*, vol. 7, 1998.

[7] M. Nikolova and R. H. Chan, "The equivalence of half-quadratic minimization and the gradient linearization iteration," *IEEE Trans Imag Proc*, vol. 16, pp. 1623–1627, 2007.

[8] J. P. Haldar, D. Hernando, S.-K. Song, and Z.-P. Liang, "Anatomically-constrained reconstruction from noisy data," *Magn Reson Med*, In Press.

[9] J. Fessler and W. Rogers, "Spatial resolution properties of penalized-likelihood image reconstruction: Space-invariant tomographs," *IEEE Trans Image Proc*, vol. 5, pp. 1346–1358, 1996.